



Delft University of Technology

Causal Inference for Spatial Data Analytics

Tomko, Martin; Xin, Yanan; Wahl, Jonas

DOI

[10.4230/DagRep.14.5.25](https://doi.org/10.4230/DagRep.14.5.25)

Publication date

2024

Document Version

Final published version

Published in

Dagstuhl Reports

Citation (APA)

Tomko, M., Xin, Y., & Wahl, J. (2024). Causal Inference for Spatial Data Analytics. *Dagstuhl Reports*, 14(5), 25-57. <https://doi.org/10.4230/DagRep.14.5.25>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Volume 14, Issue 5, May 2024

Generalization by People and Machines (Dagstuhl Seminar 24192) <i>Barbara Hammer, Filip Ilievski, Sascha Saralajew, and Frank van Harmelen</i>	1
Discrete Algorithms on Modern and Emerging Compute Infrastructure (Dagstuhl Seminar 24201) <i>Kathrin Hanauer, Uwe Naumann, Alex Pothén, and Robert Schreiber</i>	12
Causal Inference for Spatial Data Analytics (Dagstuhl Seminar 24202) <i>Martin Tomko, Yanan Xin, and Jonas Wahl</i>	25
Evaluation Perspectives of Recommender Systems: Driving Research and Education (Dagstuhl Seminar 24211) <i>Christine Bauer, Alan Said, and Eva Zangerle</i>	58
Classical-Quantum Synergies in the Theory and Practice of Quantum Error Correction (Dagstuhl Seminar 24212) <i>Carmen G. Almudéver, Leonid Pryadko, Valentin Savin, and Bane Vasic</i>	173

ISSN 2192-5283

Published online and open access by

Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany. Online available at <https://www.dagstuhl.de/dagpub/2192-5283>

Publication date

November, 2024

Bibliographic information published by the Deutsche Nationalbibliothek

The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at <https://dnb.d-nb.de>.

License

This work is licensed under a Creative Commons Attribution 4.0 International license (CC BY 4.0).



In brief, this license authorizes each and everybody to share (to copy, distribute and transmit) the work under the following conditions, without impairing or restricting the authors' moral rights:

- Attribution: The work must be attributed to its authors.

The copyright is retained by the corresponding authors.

Aims and Scope

The periodical *Dagstuhl Reports* documents the program and the results of Dagstuhl Seminars and Dagstuhl Perspectives Workshops.

In principal, for each Dagstuhl Seminar or Dagstuhl Perspectives Workshop a report is published that contains the following:

- an executive summary of the seminar program and the fundamental results,
- an overview of the talks given during the seminar (summarized as talk abstracts), and
- summaries from working groups (if applicable).

This basic framework can be extended by suitable contributions that are related to the program of the seminar, e. g. summaries from panel discussions or open problem sessions.

Editorial Board

- Elisabeth André
- Franz Baader
- Goetz Graefe
- Reiner Hähnle
- Barbara Hammer
- Lynda Hardman
- Steve Kremer
- Rupak Majumdar
- Heiko Mantel
- Lennart Martens
- Albrecht Schmidt
- Wolfgang Schröder-Preikschat
- Raimund Seidel (*Editor-in-Chief*)
- Heike Wehrheim
- Verena Wolf
- Martina Zitterbart

Editorial Office

Michael Wagner (*Managing Editor*)
Michael Didas (*Managing Editor*)
Jutka Gasiorowski (*Editorial Assistance*)
Dagmar Glaser (*Editorial Assistance*)
Thomas Schillo (*Technical Assistance*)

Contact

Schloss Dagstuhl – Leibniz-Zentrum für Informatik
Dagstuhl Reports, Editorial Office
Oktavie-Allee, 66687 Wadern, Germany
reports@dagstuhl.de
<https://www.dagstuhl.de/dagrep>

Digital Object Identifier: 10.4230/DagRep.14.5.i

Causal Inference for Spatial Data Analytics

Martin Tomko^{*1}, Yanan Xin^{*2}, and Jonas Wahl^{†3}

1 The University of Melbourne, AU. tomkom@unimelb.edu.au

2 Delft University of Technology, NL. yanan.xin@tudelft.nl

3 TU Berlin, DE. wahl@tu-berlin.de

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 24202 “Causal Inference for Spatial Data Analytics”, taking place at Schloss Dagstuhl between May 12th and 17th, 2024.

The ability to identify causal relationships in spatial data is increasingly important for designing effective policy interventions in environmental science, epidemiology, urban planning, and traffic management. Current spatial data analytic methods rely mainly on descriptive and predictive methods that lack explicit causal models. Spatial causal inference, i.e. causal inference with spatial information offers a promising tool to address this challenge by extending causal inference methodologies to spatial domains. However, this translation is challenging due to spatial effects that might violate fundamental assumptions of causal inference. Spatial causal inference is therefore still in its infancy, and there is a pressing need to accelerate its theoretical development and support its adoption with a well-grounded methodological toolset. To facilitate the necessary interdisciplinary exchange of ideas we convened the first Dagstuhl Seminar on Causal Inference for Spatial Data Analytics.

Seminar May 12–17, 2024 – <https://www.dagstuhl.de/24202>

2012 ACM Subject Classification Computing methodologies → Causal reasoning and diagnostics; Mathematics of computing → Causal networks; Theory of computation → Machine learning theory; Computing methodologies → Spatial and physical reasoning; Applied computing → Cartography

Keywords and phrases Spatial Causal Analysis, Spatial Causal Inference, Spatial Causal Discovery, Spatial Analysis, Spatial Data, Dagstuhl Seminar

Digital Object Identifier 10.4230/DagRep.14.5.25

1 Executive Summary

Martin Tomko (The University of Melbourne, AU, tomkom@unimelb.edu.au)

Yanan Xin (ETH Zürich, CH, yanxin@ethz.ch)

License © Creative Commons BY 4.0 International license
© Martin Tomko and Yanan Xin

Spatial data analytics has undergone a revolution in recent years due to the availability of large, observational spatial datasets and advances in spatially-explicit statistical analysis as well as in machine learning. Despite these improvements, the current spatial data analysis methods primarily center on exploratory, descriptive, and predictive modeling that are grounded in correlational analysis. These approaches fall short of being able to quantify (and sometimes even identify) causal relationships. However, there has been an increasing interest

* Editor / Organizer

† Editorial Assistant / Collector



Except where otherwise noted, content of this report is licensed under a Creative Commons BY 4.0 International license

Causal Inference for Spatial Data Analytics, *Dagstuhl Reports*, Vol. 14, Issue 5, pp. 25–57

Editors: Martin Tomko and Yanan Xin



Dagstuhl Reports

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

in identifying and quantifying causal relationships in spatial data which are key to designing effective policy interventions in critical applications such as environmental and population science, climate science, epidemiology, urban planning, and traffic management.

Causal inference has been an active field of study in statistics and philosophy for some time. It recently gained traction in the machine learning community as a promising method for enabling more intelligent AI capable of causal reasoning. Yet, the application of existing causal inference methods to the spatial domain is not straightforward, and a theoretical and methodological foundation for spatial causal analysis is in its infancy. Spatial effects, such as spatial dependence and spatial heterogeneity, violate the fundamental assumptions of current causal inference frameworks. Besides, the large sample size, high dimensionality (space, time, attributes), and dynamic properties of spatio-temporal data also pose great challenges in inferencing causal effects. Thus, there is a pressing need to accelerate the theoretical development in the field of spatial causal inference and enable a broader adoption of the methodological approaches supported by a well-grounded analytical toolset. Researchers in environmental sciences, spatial econometrics, spatial statistics, theoretical GIScience, and computing/machine learning communities have started making significant, yet thus far disparate efforts contributing to the foundations of spatial causal inference. This lack of interdisciplinary exchange of ideas and a comprehensive understanding of the potential applications and limitations of spatial causal inference hinders progress across these disciplines.

As machine learning rapidly penetrates various spatial decision-making processes, the time is right to enable cross-discipline conversations around spatial causal inference, and thus maximize the impact of sound methodologies. As AI becomes widely applied to spatial data analysis supporting planning and policy-making, it is imperative to develop approaches that are interpretable, grounded, robust, and responsible. Enabling the conversations between theoretical, computational, and domain experts who are active in causal inference and its application for spatio-temporal systems will accelerate the development of more intelligent and responsible AI for spatial decision-making.

This seminar is convened to initiate conversations across disciplines on these critical questions around spatial causal inference. This five-day seminar covers topics on the definitions and theories of spatial causal inference, methodologies and applications, software and benchmark datasets, and open questions. A detailed program of the seminar is provided in Figure 1. A summary of the daily discussions is shown below.

- **Unified Definitions of Spatial Causal Inference.** The discussion focused on the specification of the spatial component in the causal inference process, covering topics on the formalization of spatial causal inference questions, representations (e.g., Spatial DAG), modeling approaches, and practical relevance.
- **Methodological Challenges and Solutions.** Methodological challenges were demonstrated through case studies in environmental science, transportation, advertisement and recommendations, and other social science applications. Based on these case studies, the group explored methods and ideas for modeling spatial confounding, spatial interference, spatial treatments, and evaluation of spatial causal analysis.
- **Open-Source Software and Benchmarks.** The session featured demos of the open-source Python packages `causalml` and `tigramite`. Following the demonstrations, the group dived into the discussions of casual discovery evaluations and establishing benchmarks for spatial causal inference.
- **Open Questions and the Road Ahead.** The group proposed key research questions in the field of spatial causal inference and identified interests for continued collaborations on these topics.

	Day 1	Day 2	Day 3	Day 4	Day 5
Morning	• Welcome, programs, short introduction • Definitions and theories	• Presentations of case studies	• Continued discussions on methodologies	• Demos of open-source packages	• Summary of the seminar • Plan for post-seminar activities
Afternoon	• Spatial causal inference vs discovery • Brainstorming: questions/concepts/open problems	• Group discussions on methodologies	Hiking	• Benchmarking and evaluation • Open questions	Departure

■ **Figure 1** Program of Dagstuhl Seminar: Causal Inference for Spatial Data Analytics (May 12th – 17th, 2024).

As a major outcome of the seminar, key challenges and research questions were identified in the field, as outlined in Section 4.4.5 Open Questions and also detailed in the notes of our daily discussions. We hope these thoughts and ideas will inspire a broader research interest in spatial causal inference and continue the exchange across disciplines, as well as between academia and industry.

The seminar resulted in the desire to continue these discussions in a series of workshops (the first to take place at ACM SIGSPATIAL 2024) and the need to establish a community (`spatial-causal.org`).

In the following, the report will first present the position statements prepared by seminar participants on their thoughts related to spatial causal inference. Next, detailed notes of our daily discussions are documented in the report.

2 Table of Contents

Executive Summary

<i>Martin Tomko and Yanan Xin</i>	25
---	----

Position Statements

Introductory statement	
<i>Martin Tomko</i>	30
Introductory statement	
<i>Jianwu Wang</i>	30
Position statement	
<i>Katerina Schindlerova</i>	31
Introductory statement	
<i>Simon Dirmeier</i>	32
Introductory statement	
<i>Urmi Ninad</i>	32
Spatial statistical modelling for spatial causal inference	
<i>Andrew Zammit-Mangion</i>	33
Towards a holistic theory of spatial-causal inference	
<i>Kevin Credit</i>	34
Four open questions for a spatial causal inference	
<i>Levi John Wolf</i>	35
Emerging opportunities and challenges for spatial causal inference	
<i>Shu Yang</i>	37
Position Statement	
<i>Andreas Gerhardus</i>	37
Position statement	
<i>Jonas Wahl</i>	38
Position statement	
<i>Harinen Totte</i>	39
Position statement	
<i>Markus Reichstein</i>	39
Position statement	
<i>Cécile de Bézenac</i>	40
Position statement	
<i>Yanan Xin</i>	41
Position statement	
<i>Dominik Janzing</i>	42

Daily Summaries

Day 1: Definitions and Theories of Spatial Causal Inference	42
Day 2: Methodologies of Spatial Causal Inference	46
Day 3: Methodologies Continued	50

Day 4: Demo of Causal Inference Packages, Benchmarking, and Open Questions . 51

Day 5: Revisiting the Definition of Spatial Causal Inference 54

Participants 57

3 Position Statements

3.1 Introductory statement

Martin Tomko (The University of Melbourne, Parkville, AU, tomkom@unimelb.edu.au)

License  Creative Commons BY 4.0 International license
© Martin Tomko

Research documenting spatial causal inference is scattered across disciplines. This leads to inconsistent language describing highly heterogeneous theoretical commitments, model assumptions, data processing approaches and modelling methods. As a result, it is currently hard to synthesize best practices, unify methods under broader methodological frameworks, and provide guidance to researchers entering this nascent field. I am looking forward to the discussion unifying the perspective on spatial causal inference tasks across the disciplinary perspective represented at the seminar. In particular, I will be interested in the commitments *preceding* the data science pipeline – the translation from a theoretically grounded position that informs the design of a causal DAG; the subsequent refinement and explicit exposure of the assumed presence and role of spatial processes in the causal mechanisms captured by the DAG, and a potential additional step, that I here term *implementation DAG*, linking the capture of the causal chain to the data that will be analysed, including their fundamental properties (incl. spatial support and scale, temporal scale, measurement levels).

It is, in my eyes, necessary to overtly state theoretical positions and a grounded hypothesis before the data science pipeline for causal inference can be initiated. This subsequent pipeline (also called spatial causal framework, e.g. by [1]) needs to be grounded in such a theoretical statement, to make it clear which causal influences are analyzed and measured, and why others may be omitted.

Establishing a strong practice of overt theoretical commitments before initiating the analytical pipeline would, hopefully, support the interpretability of the studies, their replication, and the ability to judge the applicability of the results across (spatial) domains.

References

- 1 Kamal Akbari, Stephan Winter, and Martin Tomko. Spatial causality: A systematic review on spatial causal inference. *Geographical Analysis*, 55(1):56–89, 2023.

3.2 Introductory statement

Jianwu Wang (Department of Information Systems, University of Maryland, Baltimore, USA, jianwu@umbc.edu)

License  Creative Commons BY 4.0 International license
© Jianwu Wang

Spatial causal inference is still a research area in its infant stage. It deserves a lot of research. This seminar gives a great opportunity to check different opinions from attendees who are from very different disciplines/backgrounds. Some consensus was drawn from the meeting discussion. My overall position statements are:

1. **Benchmarking:** Benchmarking is critical to understanding performance differences among various solutions proposed everyday. [1] provides an overview of related causal discovery and inference applied to Earth science domain, which are mostly spatiotemporal data. It also made efforts to list related synthetic and real-world data used in related research.

2. **Machine/deep learning + Causality:** The integration of machine/deep learning with causal AI could greatly help each other. Integrating machine/deep learning into causal inference could help causal inference's performance by finding complicated patterns from data. For instance, [2] shows how deep learning can be used to estimate direct and indirect causal effects of spatiotemporal interventions in presence of spatial interference. By integrating causal discovery/inference results could help machine learning models' explainability.
3. **Taxonomy:** Some primer on the basic taxonomy/terminology will help researchers understand each other's work.
4. **Community building:** Additional community building efforts including additional rounds of Dagstuhl Seminar, workshops and tutorials would greatly help the community grow.

References

- 1 Sahara Ali, Uzma Hasan, Xingyan Li, Omar Faruque, Akila Sampath, Yiyi Huang, Md Osman Gani, and Jianwu Wang. Causality for earth science—a review on time-series and spatiotemporal causality methods. *arXiv preprint arXiv:2404.05746*, 2024.
- 2 Sahara Ali, Omar Faruque, and Jianwu Wang. Estimating direct and indirect causal effects of spatiotemporal interventions in presence of spatial interference. *arXiv preprint arXiv:2405.08174*, 2024.

3.3 Position statement

Katerina Schindlerova (University of Vienna, Währingerstrasse 29, 1090 Vienna, Austria, – katerina.schindlerova@univie.ac.at)

License  Creative Commons BY 4.0 International license
© Katerina Schindlerova

Intensive discussions among the seminar participants with the background in spatial statistics and/or causal inference have brought up ideas how to define the “spatial causal inference” as well to formulate possible objectives and working directions of this new discipline. Some of the working directions can be:

1. In case of multivariate case, it has been proposed to express the spatial distance of the variables in the neighborhood of the target as a moderation variable.
2. Predictive causal inference based in Granger [1] and its non-linear versions [2] has been an established field of causal inference, especially for temporal data/observations coming from Earth sciences. Although spatio-temporal Granger causality has been introduced for two variables [3], to our best knowledge, the extension to $p > 2$ spatio-temporal variables has not been studied yet.
3. Setting the causal discovery graph based on the dependence test (e.g. by PC [4]) and then applying causal inference to this graph (DAG) can provide ambiguous graphs for spatial scenarios, if “new” data (i.e. those with non-zero distance from the data for which was the graph generated) is used. There is a question for which types of data (or for which data distributions) this two-step procedure provides a unique output causal graph. Otherwise these types of causal discovery graphs could be used separately for data in different spatial scenarios and the inference then applied to the resulted graphs.

References

- 1 C.W.J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society*, pages 424–438, 1969.
- 2 Kateřina Hlaváčková-Schindler and Claudia Plant. Heterogeneous graphical Granger causality by minimum message length. *Entropy*, 22(12):1400, 2020.
- 3 Qiang Luo, Wenlian Lu, Wei Cheng, Pedro A Valdes-Sosa, Xiaotong Wen, Mingzhou Ding, and Jianfeng Feng. Spatio-temporal Granger causality: A new framework. *NeuroImage*, 79:241–263, 2013.
- 4 Peter Spirtes, Clark Glymour, Richard Scheines, Peter Spirtes, Clark Glymour, and Richard Scheines. Discovery algorithms without causal sufficiency. *Causation, Prediction, and Search*, pages 163–200, 1993.

3.4 Introductory statement

Simon Dirmeier (Swiss Data Science Center, ETH Zürich, CH, simon.dirmeier@sdsc.ethz.ch)

License  Creative Commons BY 4.0 International license
© Simon Dirmeier

The Dagstuhl Seminar on *causal inference for spatial data analytics* aimed to formalize where cause-effect relationships in spatial data exist, how they can be potentially discovered, and how their effect sizes can be quantified.

Briefly, let a *causal spatial inference* problem be defined as a problem involving an acyclic digraph that denotes statistical (causal) dependencies between random variables (X, Y, Z) and which has a potential spatial component s yielding the *spatial* graphical model $(X(s), Y(s), Z(s))$.

Intriguingly, in the most general case, the factorization of the joint distribution defined by the spatial graphical model contains both causal dependencies, e.g., $X(s) \rightarrow Y(s)$ as well as statistical ones, e.g. encoded via potential functions $\psi(X(s), X(t))$, which can complicate typical operations, such as computation of interventional or counterfactual distributions. In specific cases, e.g., when no correlation structure between spatial variables exists, the spatial causal inference problem seems to be readily reduceable to the conventional causal framework which allows for structure learning with contemporary constraint-based methods and effect estimation using, e.g., the Pearllean identification criteria.

The emerging field of causal spatial inference offers a wide variety of interesting future research directions ranging from discovery to estimation.

3.5 Introductory statement

Urmi Ninad (Causal Inference and Climate Informatics Lab, Technische Universität Berlin, DE, urmi.ninad@tu-berlin.de)

License  Creative Commons BY 4.0 International license
© Urmi Ninad

Causal inference aims to formalise the investigative query of discovering and quantifying pathways of causation between variables. This query arises within many sciences, such as climate science, neuroscience and geography. In several cases of interest, the variables, or

the interactions between them, or the causal structure itself varies over space. A closer look at the several applications cases quickly illustrates the richness of the problem of incorporating spatial statistics into traditional causal inference language. We studied in detail the problem of quantifying the effect of CO_2 on gross primary production (GPP) of plants. We also investigated disentangling the causal effect of emissions from a power plant at certain space-coordinate from the causal effect of derivative (also called *children* in causal graphical language) variables around that space-time coordinate also presents several challenges, starting with the complication that the two causal drivers of air quality are non-trivially related to, and influenced by each other. In the course of discussing these and a few other examples, novel problems emerge, for which the causal inference toolbox is found wanting.

In this Dagstuhl Seminar, a fruitful exchange between the spatial statistics and the causal inference community resulted in the perspective that “spatial causal inference” is, in fact, an umbrella term for problems where the space as a dimension plays a role, and either cannot be ignored to ensure soundness, or can be instrumental for certain computations, such as that of de-confounding. In order for this field to progress, a multi-pronged approach is required that is motivated by grouping spatial causal inference tasks into clusters and advancing them individually. The task of establishing a unified framework for any and all causal inference queries that use the space dimension non-trivially would ideally emerge thereafter.

3.6 Spatial statistical modelling for spatial causal inference

Andrew Zammit-Mangion (University of Wollongong, AU, azm@uow.edu.au)

Spatial causality can express itself in various ways; it is not straightforward to represent in directed acyclic graphs, and special care must be taken when establishing equations and governing notation. The following insights from spatial statistical modelling may be useful to bear in mind when constructing spatial causal models:

1. **Causality is a property of the underlying process:** We should resist the temptation to think of a spatial treatment as directly affecting the outcome (or observation); the spatial treatment affects a spatial variable that may only be observed through incomplete and noisy data. Causality is between the spatial treatment and an (unobserved/latent) spatial variable. This distinction is somewhat critical, because
2. **What you see is not what you want to get (a.k.a. think hierarchically) [1]:** Even if the spatial outcome is observed in its entirety (i.e., there are no missing values), it is generally a noisy version of the underlying process. Interest is in the causal effect on the process, and not on the noisy data. Measurement errors, biases, etc. need to be factored in when making inference on the causal effect.
3. **Think continuously:** In spatial causal models, treatments may be point referenced, and data might be areal, or vice versa; this “change of support” problem [2] can be solved by modelling everything on a continuously-indexed spatial domain, and by defining treatments, outcomes, and any confounding variables as integrals over the spatial domain; see also [3, 4].
4. **Treatments can have far-reaching consequences:** A treatment at s_i may cause the outcome to change somewhere far away; at s_j say. This is clearly the case in environmental problems where polluting rivers affect ecosystems downstream, or where toxic gases from a chemical plant affect people living in villages downwind. One therefore often needs to model the “sensitivity” of the outcome at s_j to a treatment at s_i , potentially for

every $\mathbf{s}_i, \mathbf{s}_j$ inside the spatial domain of interest. It is from the combination of this sensitivity and the treatment footprint (e.g., their inner product) that the effect on $\mathbf{s}_j$ can be established.

5. **Think temporally:** Spatial variables are either temporal snapshots of spatio-temporal variables or averages of spatio-temporal variables over time. A legitimate question is: Does averaged spatio-temporal causality lead to spatial causality? I believe the answer to this question is yes (under linearity and some other assumptions): if the outcome at a certain point in space and time is caused by a convolution of the spatio-temporal treatment and a spatio-temporally varying sensitivity, then the temporally-averaged outcome at that spatial location can be obtained from the the temporally averaged treatment and the temporally aggregated sensitivity. This is a consequence of Fubini’s theorem; see [5], Section 4.3, for details.

References

- 1 Noel Cressie and Christopher K Wikle. *Statistics for Spatio-Temporal Data*. John Wiley & Sons, 2011.
- 2 Christopher K Wikle and L Mark Berliner. Combining information across spatial scales. *Technometrics*, 47(1):80–91, 2005.
- 3 Matthew Sainsbury-Dale, Andrew Zammit-Mangion, and Noel Cressie. Modeling big, heterogeneous, non-Gaussian spatial and spatio-temporal data using FRK. *Journal of Statistical Software*, 108:1–39, 2024.
- 4 Andrew Zammit-Mangion and Noel Cressie. FRK: an R package for spatial and spatio-temporal prediction with large datasets. *Journal of Statistical Software*, 98:1–48, 2021.
- 5 Noel Cressie and Andrew Zammit-Mangion. Multivariate spatial covariance models: a conditional approach. *Biometrika*, 103(4):915–935, 2016.

3.7 Towards a holistic theory of spatial-causal inference

Kevin Credit (Maynooth University, IE, kevin.credit@mu.ie)

License  Creative Commons BY 4.0 International license
© Kevin Credit

Causal inference is an important approach for providing useful answers to scientific questions – and solutions to applied problems – in regional science, geography, and urban planning. However, there are a number of challenges to using causal inference in urban-geographic settings: 1) the overlapping correlations inherent in spatial data and data-generating processes often violate the basic assumptions of the potential outcomes model; 2) many of the causal effects of interest in these settings are spatially- and temporally-heterogeneous, and adoption/treatment is often staggered and/or of varying intensity; 3) in many cases true randomized experiments in these settings are not possible to design to answer the research questions of interest, which makes the availability of appropriate secondary data and methodological choices of individual researchers particularly important.

Beyond these challenges, it is also important to note that the theory and methods of causal inference have developed in somewhat distinct literatures and are often used on different kinds of data to answer different kinds of questions. For instance, Rubin’s potential outcomes model originated in statistics [1] and is used in a wide range of domains, often in the social and health sciences. The difference-in-differences method, which was applied seminally in Card and Krueger’s analysis [2] of minimum wage and employment – for which

Card was awarded the 2021 Nobel Prize in Economics – is by far the dominant framework for causal inference in economics, and continues to influence approaches to causal inference coming from the spatial econometric perspective using (typically) areal spatial data [3]. Other methods of causal inference, such as the Structural Causal Model (SCM) [4], are used to study earth systems data and spatial-temporal processes. Even more recently, new approaches for estimating heterogeneous treatment effects – such as the “metalearners” [5] and causal forest [6] – have emerged from the machine learning literature.

While the underlying philosophy of these approaches are arguably the same, they currently speak different languages, use different notation, and focus on different assumptions. Thus I think that any “spatial-causal” project must first acknowledge the unique development of the various strands of causal inference in different disciplines. It should also attempt to build a more holistic theory – or at least a more holistic accounting – of causal inference as applied to spatial data, starting from general principles and moving to more specific assumptions and approaches that can be applied to different kinds of spatial data in different substantive domains. In my view, spatial causal inference should be a “broad church” that includes any work dealing with problems of causal inference where the underlying process of causation – including treatment, outcome, susceptibility, or confounding – varies in space and is accounted for somehow in the analysis.

References

- 1 Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688, 1974.
- 2 David Card and Alan B Krueger. Minimum wages and employment: a case study of the fast-food industry in new jersey and pennsylvania: reply. *American Economic Review*, 90(5):1397–1420, 2000.
- 3 Marynia Kolak and Luc Anselin. A spatial perspective on the econometrics of program evaluation. *International Regional Science Review*, 43(1-2):128–153, 2020.
- 4 Judea Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688, 1995.
- 5 Sören R. Künzel, Jasjeet S. Sekhon, Peter J. Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10):4156–4165, 2019.
- 6 Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018.

3.8 Four open questions for a spatial causal inference

Levi John Wolf (University of Bristol, United Kingdom – levi.john.wolf@bristol.ac.uk)

License  Creative Commons BY 4.0 International license
© Levi John Wolf

Geographic information science has a causality problem – for too long, it has focused on defining causality as a kind of regularity, rather than as something arising from difference-making interventions [1]. This has made it challenging to think about a very wide variety of important geographical and urban planning problems, such as the expected effect of opening transit stations, instituting new governmental policies, or intervening on the natural (or built) environment to improve ecosystems and the environment. Despite many attempts to push the field into an intervention-focused framing, GIScientists have remained focused on laws due to the prevailing fixation on idiographic-nomothetic debates in 20th century geography

reverberating through contemporary discussions of reproducibility and generalizability. It is time to adopt more useful theoretical frameworks and mathematical formalisms in which geographical planning, policy, and intervention can be understood [2]. This foundation for a more causally-oriented geographical analysis requires a few important and fundamental innovations in *spatial causal inference*.

Spatial DAGS It is quite challenging to understand how to appropriately represent contextual effects in directed acyclic graph (DAG) representations of models. These are the bread and butter of contemporary causal inference, yet it is challenging to understand how to represent spatial concepts within them. After this seminar, it seems that one useful way forward may be through chain graph concepts, which require us to specify *process-specific forms of spatial dependence* within the DAG itself.

Attributing Spatial Context Contextual effects are a very important component of spatial planning and program evaluation. Broadly speaking, this refers to the effect that surrounding conditions have on an outcome. In an interventionist case, it refers to the effect that surrounding *treatment* may have on your outcome. Distinct from spatial interference (where surrounding treatment *interferes with* your treatment), this is an important novel component for spatial causal analysis which is difficult to represent in classical causal analytical frameworks. This is distinct from the much more difficult example of spatial endogeneity within Y , as might happen when outcomes that are near one another influence one another.

General Spatial Causal Model Placing contextual effects alongside other well-studied spatial causal issues (such as spatial confounding or interference), it becomes important to define a so-called *general spatial causal model* that can be used to simulate data according to multiple different processes. It seems that only by combining these various processes can we actually identify treatment effect estimators that are robust to these processes.

Spatial Targeting In classical causal inference, I have learned this week that *targeting* is the practice of identifying which subset would most benefit from treatment. Classical targeting approaches assume that individuals' treatments can be administered independently, but this is not so when treatment has spatial components. In a *spatial targeting* problem, applying a treatment in a given location (or with a given distance decay effect) may not be beneficial to the surroundings of the treatment, even though it is beneficial at the site of treatment. In GIScience, it is important to try to identify *where* an intervention might maximize the post-intervention difference. This has mathematical similarities to a maximal covering location problem (MCLP), a kind of set covering problem, where it would be important to identify specific target locations that maximize benefits as a function both of the direct and spillover treatment. Solving this remains an open question.

Regardless, this seminar has been quite effective in stimulating cross-community collaboration between spatial statisticians and computer scientists studying spatial causal analysis (or any subset of those terms). I believe this was immensely valuable, and it will undoubtedly influence my thinking and future work.

References

- 1 Jing Zhang and Levi John Wolf. Rethinking “causality” in quantitative human geography. *Geography Compass*, 18(3):e12743, 2024.
- 2 Levi John Wolf. Confounded Local Inference: Extending local Moran statistics to handle confounding. pages 1–16.

3.9 Emerging opportunities and challenges for spatial causal inference

Shu Yang (North Carolina State University – Raleigh, USA, syang24@ncsu.edu)

License © Creative Commons BY 4.0 International license
© Shu Yang

Spatial causal inference focuses on estimating the effect of treatments, interventions, exposures, or policies, and inferring causal relationships using spatial data. The Dagstuhl Seminar on spatial causal inference is both timely and important for fostering discussions among experts from diverse fields and backgrounds.

The importance of spatial causal inference is underscored by the emergence of numerous scientific questions that are inherently causal in nature [1], and the increasing availability of large studies containing spatial data, such as those in environmental health, epidemiology, geoscience, economics, urban planning, and earth science. Despite its potential, spatial causal inference remains in its early stages and encounters significant challenges.

A fundamental characteristic of spatial data is that variables located closer together tend to be more similar than those further apart, as per Tobler’s First Law of Geography. This spatial correlation can violate classic causal assumptions, such as the independence and identical distribution of observations, and the stable unit treatment value assumption, where an outcome at one location may depend not only on the treatment at that location but also on treatments at nearby locations. Additionally, causal relationships may vary spatially due to differing environmental conditions across large areas. Addressing these complexities in spatial causal inference is challenging but essential.

Nevertheless, spatially structured data can be an asset rather than a drawback. The inherent structure can be leveraged to enhance causal inference. For example, different spatial patterns in outcomes and confounders can be utilized to mitigate biases resulting from missing spatial confounders [2]. This potential makes me optimistic about the future of spatial causal inference.

References

- 1 Brian J. Reich, Shu Yang, Yawen Guan, Andrew B. Giffin, Matthew J. Miller, and Ana Rappold. A review of spatial causal inference methods for environmental and epidemiological applications. *International Statistical Review*, 89(3):605–634, 2021.
- 2 Yawen Guan, Garritt L Page, Brian J Reich, Massimo Ventrucci, and Shu Yang. Spectral adjustment for spatial confounding. *Biometrika*, 110(3):699–719, 12 2022.

3.10 Position Statement

Andreas Gerhardus (DLR Institute for Data Science, Jena, DE, andreas.gerhardus@dlr.de)

License © Creative Commons BY 4.0 International license
© Andreas Gerhardus

In my research, I work on theory and methods for causal inference as well as the application of these methods to real-world data. My main focus lies on time series data, but I am also increasingly dealing with applications to spatio-temporal data. This is why I am enthusiastic to participate in and contribute to a week of discussions on causal analysis for spatial data. In my view, it is particularly important to discuss the importance of specifying the respective targets of estimation and discovery before proceeding to derive estimands and estimates for these targets. Moreover, I am looking forward to discussions on the different roles that

space can play in the analysis. For example, in one case one might be interested in the cause-and-effect relationships between variables that by themselves are spatial, whereas in another case one might be interested in the relationships between the variables at individual space points and to this end need to take care of the confounding effect by variables at the other spatial locations. From the seminar, I hope to take away thoughts for future work and to lay the foundation for potential collaborations in the future.

3.11 Position statement

Jonas Wahl (Technical University of Berlin, DE & DLR Institute for Data Science Jena, DE, wahl@tu-berlin.de)

License  Creative Commons BY 4.0 International license
© Jonas Wahl

As many examples brought forward during this seminar aptly demonstrated, many questions on cause and effect in the sciences involve data that is inherently spatial. On the one hand, geoscientists, econometricians and statisticians have developed practical techniques to deal with spatially structured data (many of which were discussed and presented at this seminar), while causal inference researchers have formalized the notion of causal effects and interventions with a focus on non-spatial data. In my opinion, there is a need for models that clearly delineate between causally induced and spatially induced relationships, and on which the concept of an intervention is unambiguously defined. These models should be close to actual scientific practice. Therefore, equipping existing models with a notion of intervention that implies a definition of causal effect which matches applied researchers' intuition would be a useful step forward. Hierarchical spatial process model [1] seem to be a particularly fitting model class for this goal as they enable to formulate causal notions on the level of the underlying process instead of directly on the measured data. Another reference that has crystallized as a useful starting point for incorporating the causal notion explicitly into spatio-temporal models is [2]. In addition, tools for generating data with both causal and spatial components and literature that reviews existing ideas in combination with practical examples and code would help the community significantly.

The Dagstuhl Seminar has done a particularly great job in making researchers explain their methodology to experts from other fields, spurning interactions that would have been rare otherwise. To keep the community going, future meetings, whether at Dagstuhl or in other venues, would be crucial and first steps have been taken towards that goal.

References

- 1 Noel Cressie and Christopher K Wikle. *Statistics for Spatio-Temporal Data*. John Wiley & Sons, 2011.
- 2 Steffen L. Lauritzen and Thomas S. Richardson. Chain Graph Models and their Causal Interpretations. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 64(3):321–348, August 2002.

3.12 Position statement

Totte Harinen (Airbnb – San Francisco, USA, totte.harinen@airbnb.com)

License  Creative Commons BY 4.0 International license
© Harinen Totte

Spatial causal inference has the potential to be highly relevant for data scientists working in industry. One immediate application is A/B testing, where spatial information can be used to understand treatment effect heterogeneity and the regional targeting of interventions. Ideas discussed in the seminar include spatial versions of existing causal machine learning algorithms and using spatial features as covariates. Spatial causal inference can also help with more well-known problems such as spatial confounding and spillover effects.

Industry can also benefit from the conceptual ideas discussed in the seminar, including ways to reason about and represent space in the context of inference. Because there is no systematic ways to represent space in data science, assumptions about its influence are often left unstated. Developing frameworks for spatial causal reasoning would therefore plausibly improve decision-making in industry.

To move spatial causal inference forward in industry, we need case studies that show its successful application in concrete problems with non-curated data. Working on such case studies will likely also surface interesting new methodological challenges.

3.13 Position statement

Markus Reichstein (Max-Planck-Institute for Biogeochemistry & ELLIS Unit Jena, DE, mreichstein@bgc-jena.mpg.de)

License  Creative Commons BY 4.0 International license
© Markus Reichstein

Both spatial dependence and causality are still often ignored in Earth system data analysis and machine learning, leading to potentially biased results and misleading conclusions about environmental effects on vegetation and ecosystems. That's why this seminar on Causal inference for Spatial Data Analysis is double-important.

Spatial dependence can be curse or blessing and I find it most interesting to see how it can help with adjusting for confounding effects. For instance if we want to quantify the effect of CO₂ on vegetation photosynthesis from observations, we have to consider confounders such as nitrogen deposition and should exploit those have a different spatial structure than CO₂ concentrations.

Another question is how to identify causal effects in a hybrid modelling framework, which combines a process-based model with a machine-learning approach.

Other ideas inspired from the seminar:

- Can we identify temporally or spatially varying DAGs from a spatio-temporal data set?
- Can we use administrative boundaries to adjust for human confounders when we want to find effects of climate variables?
- How can we identify causal spatial context effects (aka convolutions in machine learning), for instance on ecosystem responses to drought stress?

3.14 Position statement

Cécile de Bézenac (University of Leeds, GB & The Alan Turing Institute – London, GB, cdebezenac@turing.ac.uk)

License  Creative Commons BY 4.0 International license
© Cécile de Bézenac

A major endeavour of the social sciences is to explain social process and therefore to generate causal rather than associational claims. However, the identification of causal effects can prove to be challenging if not impossible in the presence of social and spatial complexity. Namely, interaction (direct or indirect) between individuals or events within a geographic environment can result in spatial mechanisms that is difficult to completely disentangle in the form of clear directed causal relationships. Having acknowledged such dynamics, what is essentially put into question is the validity of the causal claims that stem from spatial prediction models as well as from causal models that do not explicitly address the spatial nature, or the “spatiality” of the problem. From this observation emerges the importance but also the challenges of *spatial causal inference*. Defining this term and describing the holistic landscape of this burgeoning field have been the focal points of this seminar. If one were to think of spatial causality, as we have, as the more general form of causality, then one may also consider what it takes to cement spatial considerations in the causal framework. From the very interesting discussions I have noted several aspects that I see as “directions” for this:

- Methods for identification (or falsification) of “spatial effects” and learning relevant spatial representations: this problem relates to one of invariance search under spatial representations and transformations (permutation, aggregation...) for the identification (or falsification) of “spatial effects” in the problem.
- An appropriate formalism for the problem, as both a tool to reflect on the role of space and to describe it. A formalisation should translate the relevant spatial (and temporal) situation of the problem while being as actionable as possible. In the case of a DAG-based descriptions, a spatiotemporal process may blur the distinctions between nodes and edges (or what happens between two nodes). Solutions relate to the embedding of spatial nodes in the graph or to the embedding of nodes in space (implying often a change of scale or change of support). An interesting prospect would be to also consider the “embedding” of edges in space: how do spatial relations relate to causal ones? (ex: perhaps considering the use of chain graphs)
- Inference methods: One of the objectives of developing community around this topic is to bring together the tools that have been developed in various strands of research. This also requires understanding their differences in order to build a structured set of methods. I am interested in how representations of space and spatially formalising the causal problem can support the modelling choices (ex: how may a causal forest method integrate spatial information?)
- Evaluation methods: In the absence of empirical ground truth, I believe one way forward is to integrate our assumptions on the underlying spatial processes in a simulation framework. By distinguishing between types of spatial (and temporal) mechanism one can combine to generate a space of assumptions that would serve as a set (collection of artificial datasets or a simulation protocol) on which to tests spatial causal methods.

I am interesting in exploring the synergies between empirical and simulation-based methods in the context of spatial causal inference, particularly under complexity assumptions. In that perspective, a complex systems approach such as Agent Based Models could serve the

development of statistical tests and inference methods. Furthermore, I believe there is ground for developing multi-agent simulation-based inference methods for spatial causal analysis, harnessing their ability to describe the transition from micro-level processes to emergent properties. On the condition of course that the questions of uncertainty and validation be systematically addressed.

3.15 Position statement

Yanan Xin (*Institute of Cartography and Geoinformation, ETH Zürich, CH, yanxin@ethz.ch*)

License  Creative Commons BY 4.0 International license
© Yanan Xin

There has been a growing research interest in reasoning causality in spatial data analysis, aiming to answer “what if” questions beyond “what is” queries. A fundamental challenge in spatial causal inference is to understand, identify, and quantify the influence space exerts in causal inference processes.

This research direction involves two key components. On the theoretical side, we need to develop fundamental analytical frameworks for spatial causal inference, for example, extending the potential outcome framework or the structural causal models to explicitly account for the spatial dimension. On the methodological side, such as data-driven causal discovery and causal effect estimation, we need robust approaches to distill causal relationships from data. It is also important to understand in what situations these causal relationships will be infeasible to identify or quantify due to spatial interference and spatial confounding.

The spatial factor also poses challenges to causal machine learning research [1]. Previous studies have highlighted that integrating causality can enhance the interpretability and robustness of machine learning models. This is particularly appealing for spatial data science, where the lack of transparency and generalizability of machine learning models hinders their adoption in various spatial applications to support decision-making [2]. In recent years, some causal machine learning approaches have been developed in this direction [3]. However, accounting for the spatial dimension remains challenging. For example, in causal representation learning or causal generative modeling, the spatial dimension can either be considered as a factor to group other features or it can be considered as a separate feature. How the formulation influences the analysis and interpretation of results deserves further investigation.

References

- 1 Jean Kaddour, Aengus Lynch, Qi Liu, Matt J Kusner, and Ricardo Silva. Causal machine learning: A survey and open problems. *arXiv preprint arXiv:2206.15475*, 2022.
- 2 Yanan Xin, Natasa Tagasovska, Fernando Perez-Cruz, and Martin Raubal. Vision paper: causal inference for interpretable and robust machine learning in mobility analysis. In *Proceedings of the 30th International Conference on Advances in Geographic Information Systems*, pages 1–4, 2022.
- 3 Rushan Wang, Yanan Xin, Yatao Zhang, Fernando Perez-Cruz, and Martin Raubal. Counterfactual explanations for deep learning-based traffic forecasting. *arXiv preprint arXiv:2405.00456*, 2024.

3.16 Position statement

Dominik Janzing (Amazon Web Services, Tübingen, DE, janzind@amazon.com)

License  Creative Commons BY 4.0 International license
© Dominik Janzing

How should we evaluate all these *causal discovery* methods out there? Datasets with reliable ground truth are rare, and evaluation on simulated data are questionable [1]. To enable significant progress in the field we need extensive benchmarking rather than discussing *plausibility* of results for a few datasets. To this end, we need a new theory of falsification of causal explanations that tells us the testable implications entailed by a causal explanation, following the spirit of Popper’s theory of science. To this end, we are working on “self-compatibility” and test whether outputs of causal discovery algorithms are compatible across different subsets of variables [2]. This way, algorithms can be falsified without ground truth. Although not contradicting itself is not a guarantee for being true, we have discussed notions of compatibility that are so strong that they can at least provide some evidence for the output of the algorithms. However, one of the most challenging questions raised by this approach is the definition of a good measure of compatibility together with a notion of calibration that tells us whether the observed inconsistencies are “many” or “few”. After all, perfect compatibility will never be achieved – how do we decide whether the number of contradictions is small enough to ensure that the algorithms are still useful? This is ongoing work!

References

- 1 Alexander Reisach, Christof Seiler, and Sebastian Weichwald. Beware of the simulated dag! causal discovery benchmarks may be easy to game. NeurIPS 2021.
- 2 Philipp M. Faller, Leena Chennuru Vankadara, Atalanti A. Mastakouri, Francesco Locatello, and Dominik Janzing. Self-compatibility: Evaluating causal discovery without ground truth. AISTATS 2024.

4 Daily Summaries

4.1 Day 1: Definitions and Theories of Spatial Causal Inference

On day one, we dived into the definitions and theories of spatial causal analysis, addressing fundamental questions – What is *spatial*, *causal*, and *spatial causal*?

4.1.1 Spatial causal analysis: definition and theories

Note taker: Shu Yang, edited by Yanan Xin

In this session, we commenced the discussion on the *definition* of spatial causal analysis. The participants set off by a refresher on the *do* calculus notation, through an illustration. Questions have been raised about the interpretation of causal relationships as *context-free* or *context-bound*. Note that a casual relationship is invariant to context, as illustrated by the *do* calculation. We further discussed the formulation of **Spatial Causal DAG**. The discussion centered around the following questions: When does location context matter in this formulation, what does it mean, and how does it relate to DAG? Is location just an index in the formulation of spatial causal DAG or should location be encoded as a *s* node in the DAG? How should we encode the directions between locations? Time is directional,

could we draw inspiration from encoding temporal dynamics to represent spatial directions in spatial causal DAG? How do we represent spatial interference in spatial causal DAG? Some ideas were proposed to address these questions, such as: 1) denoting the node of the spatial causal DAG as time- and site-specific. 2) using stacked DAG which offers the flexibility to easily represent spatial confounding and spatial interference concepts.

Next, we discussed in what situations space should be considered as a causal factor. Some example scenarios are given, such as when connected locations cause treatment interference or treatment-outcome relationship is geo-related. To help us better define spatial causal DAG, we discussed its potential connections with causal models defined in the iid settings and temporal causal models. Causal models built for iid data can be considered as a subset of spatial causal models, representing a special case of causal inference in which space and time do not matter. In time-series causal discovery (without considering space), the sliding window approach is used to create iid data. Similar ideas can be applied to spatial problems. For example, snapshots of space can be used to generate iid data in space. These ensembles of snapshots or grid cells are independent of each other, however, this transformation also results in information loss.

Another perspective that space only matters in the causal inference process if one has omitted critical spatial variables was also voiced, but this was argued to not hold in cases where the spatial spillover effect exists or space serves as a proxy variable.

The discussion moved on to defining the types of causal questions that are of interest to Spatial Causal Inference. Ultimately we want to answer why a causal question is spatial. One idea that emerged from the discussion is to look at the underlying *spatial process*, not just the DAG. Variables change in space or the causal relationships change over space – these relate to mechanisms of spatial processes.

4.1.2 Spatial causal analysis: inference and discovery

What is the difference between inference and discovery? The current mainstream definition of Causal Inference encompasses discovering causal relationships (causal discovery) and causal effect quantification (causal inference). Often these two aspects are intertwined. Take the PC algorithm for example. The first step is to learn the DAG structure and the second step is to check the strength of the causal relationship (effect quantification). Machine learning can be used for effect quantification and inference. For example, causal representation learning uses high-dimensional observations to learn low-rank latent variables/latent models. These latent variables and models are used to approximate causal variables and causal relationships. Another line of research is abstraction learning which uses the latent variables to generate high-dimensional observations as a way to learn simplified causal relationships. A couple of challenges exist in causal representation learning, e.g., how to account for unmeasured confounders or unmeasured causes? can causal representation learning identify the missing variables? is the causal relationship or the DAG unique? Different DAGs might be compatible in some aspects. Two DAGs might be indistinguishable, given the observations and/or assumptions. We also discussed the definition of spatial confounding. Two definitions emerged: Def 1 – neighbor confounder affects your exposure and outcome. Def 2 – a confounder that has a spatial structure.

4.1.3 Brainstorming: Questions/Concepts/Open Problems in Spatial Causal Inference

Note takers: Katerina Schindlerova, Levi Wolf, and Yanan Xin

In this session, seminar participants are asked to brainstorm questions, concepts, and open problems in spatial causal inference. These ideas are grouped into topic categories and summarized below.

Modelling.

- Model learning: learning partial weights to generate estimand – non-parametric
- Model estimation: weights = estimand – parametric
- Spatial matching

Diagnostics – How to know if you should use a spatial model.

- If the only “spatial” component is my treatment, can I just do a standard causal analysis?
- In which application we can reliably exclude spatial influence and in which not?
- A causal question, when incorporating spatial relationships or locations improves our understanding of the mechanism and size of the effect.
- Should average/individual treatment effect (ATE/ITE) be extended to add spatial and/or temporal parameters, e.g., lagged ATE?
- Clearly define a) spatial confounding b) spatial interference c) spatially varying effects d) spatial correlation; Is it possible to differentiate them from data?
- How to test that space causally matters, i.e., $P(Y|S, X) = P(Y|X)$? i.e. Y is conditionally independent of S given X.
- In a concrete causal spatial model, should the spatial variable be a conditional variable or a direct variable?
- Is the inclusion of space variable required for causal sufficiency?
- When does spatial causal analysis = classical causal analysis?
- Can we use existing techniques like chain graph models to distinguish between “spatial correlations” and correlations induced by causation?
- Inference (spatial) accepting treatment in area t_i affects all nearby treatments t (i from η)
- Spatial assignment: treatment t_i depends on spatially dependent variables z_i that also affect y_i directly
- Policy spillover: adopting treatment t_i in a place affects nearby outcomes y_i , i from η
- Spatial mediation: treatment affects outcome through a spatially varying mediator (must be distinguished from confounding!)
- Price signaling (endogeneity): y_i affects y (i from η) nearby
- Exogeneous spillover: nearby exogenous conditions X_i affect outcomes y_i (i from η) nearby regardless of treatment

Philosophical.

- Does cause always precede effect?
- Does cause have to always happen?
- How do current spatial statistics answer questions of cause and effect? Or, do they refrain from that and focus on prediction only?

Applications/Use Cases.

- What is the effect of a change in short-term rental regulation in a given city?
- What is the effect of a billboard campaign on sales?
- What is the effect of building a public transit line on CO₂ emissions?
- How can I estimate the effect of an intervention that causes customers to purchase a product when there's a limited quantity of it?
- Inter-gene signaling pathways between cells (spatial component)
- Infer causal effect of political policies between neighboring states
- What is the induced demand of a new service station?
- Does household/city/area size have an impact on resilience/decay/growth?
- Impact of lockdown on mental health?
- Peer effects in PV adoption likelihood/rate?
- Role of geographic/strategic position on climate agreement participation? (spatial relations, like network position or climate differences?)

Shared Datasets, Generators, and Benchmarks.

- Cross validation vs (and?) crossvalidation
- AB testing in spatial contexts
- Prediction uncertainty and estimate coverage (like scoring)
- How do we actually benchmark causal estimates?

Evaluation Metrics and Reported Characteristics. *(How to write evaluations of spatial causal analysis?)*


- What are the good metrics to evaluate spatial causal discovery/inference results?

Causal Discovery.


- Sample question – Which location's variables are the causes of an effect?
- Combine physical modeling to aid causal discovery?
- Can one identify spatially varying DAGs?
- Causal representation learning with spatial data

Model Formalisation (DAG++).

- In proposing a spatio-causal model, how should space be functionally expressed?
- How to define “close” and “distant” in a causal-spatial graph?
- What is a proper structure to model spatial causal relationships?
- Good ways to present assumptions? Causal DAG or ignorability assumptions in the PO framework?
- Defining spatial causal analysis in the non-ensemble setup: define it as a spatial stationarity binding or non-binding problem. That is, if the causal graph G remains “stationary” across the stack and there is no spatial autocorrelation, then it is NOT a spatial causal problem.
- How to define the vicinity of potentially causal variables in a spatial causal graph?
- Defining spatial causal analysis from the lens of target graphs. (as stacks thereof)
- How does “do” calculus change when X , Y , and Z are correlated vectors?
- Spatial causal analysis concerns itself with discovery and effect estimation in models of the following form, where DAGs are correlated over space.



- How to model heterogeneous spatial and temporal scales/granularity in causal graphs?
- Two possible ways of representing space in a DAG:



Discipline Definition (Spatial Causality).

- When space is related to treatment, effect, or confounding intensity (or direction)?
- Some variable of interest is spatial.
- Dependencies change in space if of interest or not...
- Even if the question itself is not spatial, you have spatial data to answer it
- Variables of interest include some measured variables of different locations
- Spatial relationships are strong causal links for the target estimate
- Is one of the biggest advantages of spatial causal analysis that we can “improve influence”/“borrow strength” and quantify causal relationships that would be impossible otherwise?

4.2 Day 2: Methodologies of Spatial Causal Inference

On day two, seminar participants presented case studies, showcasing methodological challenges in causal inference across different spatiotemporal applications. Based on these case studies, we grouped similar topics and continued the discussion in two separate working groups. Here we summarize the case studies and highlight key insights from the group discussions.

4.2.1 Case Study Presentations

Jianwu Wang

- **Case study on quantifying the causal impact of climate variables on arctic sea ice loss by sub-regions**
- goal is to estimate the direct, indirect, and lagged treatment effects under both temporal confounding and spatial interference with an estimation strategy

Kevin Credit

- **What are the impacts of building a new pedestrian/cycling infrastructure on adjacent residential construction, on-road CO2, and retail activity?**
- Test case: the 606 elevated trail in Chicago
- Goal is to estimate the strength of potential heterogeneous effects
- Treated area is compared to both a “close but untreated” and a “distant but still in Chicago” controls.

Katerina Schindlerova

- **Wind farm productivity across wind turbines:** Detection of climatological variables by multivariate Granger causality having a temporal influence on extreme and moderate windspeed of each turbine;
- The question is how to integrate spatial information of each turbine into one causal model for the whole farm.

- **Chicago crime count dataset:** Granger multivariate causal model for variables following Poisson distributions was used for count time series representing the numbers of daily committed crimes in Chicago and the temporal influences among various types of crimes. The question is how to integrate both temporal and spatial influences among various committed crimes into one model; More specifically, the question is to propose a plausible indexing of spatial and temporal proximity.

Shu Yang

- **Wildfire effects on air quality;**
- Effects measured by pm2.5, using propensity score matching to associate observations

Markus Reichstein

- **CO2 affects general primary productivity (GPP)**
- GPP is affected by CO2, but also meteorological features and nitrogen deposition
- An unknown “spatial process” might be omitted as a confounder that affects the ability to infer the other effects.
- **Predicting greenness outcome over summer using landscape and geographic factors:**
- given a baseline model for “standard” predictions, effective predictions can be made for GPP;
- Then there may be some very local factors geographically (aspect, local hydrology) that may be relevant

Totte Harinen

- **AirBnB ranking experiments:**
- First, change the ranking of the Airbnbs that are presented according to a search, then compare the average booking (frequency?) between treated and controls.
- Another variant would be to “increase” the suitability/ranking of Airbnb listings based on their quality. This makes some listings appear exclusively in some searches, creating an exclusion-based interference on the exposure to a listing itself (you book what they show you, and assignment into treatment/control means you get a distinct set of properties)

Levi John Wolf

- **Geographic regression discontinuity (NYC house prices)** <https://doi.org/10.1080/01621459.2020.1817749>, how do school districts affect house prices?
- Estimand is the premium for being within a school district, based on the sale price of a house

Martin Tomko

- **SatNaV supported Navigation with gaps (work with Kamal Akbari):**
- People navigate from POIs and home locations. Frequently, they get navigation information to get to destinations. But, they turn off the navigation at particular times. What drives the factors behind how people induce gaps? What causes them to turn it off? The theory is that environmental familiarity confounded with other exogenous factors about the neighborhood is at play;
- **Train station opening effect on house prices (work with Kamal Akbari),**
- RDD, treatment as the opening of the station,
- treated/untreated units may also “spill over” across different adjacent stations (multiple coverages?)
- **Malaria and outdoor movement behavior (work with Buran Cong and Wila Wu):**

- How likely are you to be infected with malaria after spending time in a forest (in Cambodia)?
- Exploring the role of trajectory sampling on inference. Fine-grained trajectories can lead to additional potential complexity in how the exposure is modeled, and many different local factors (e.g. water features or forest fragmentation) can also modify the exposure to malaria-causing factors, coarse trajectories neglect nuances in the exposure.

4.2.2 Discussion about causal discovery and thematic grouping

The session continued by clustering the research directions by topics/applications, but also by nature of the tasks:

- Discovery problems relating to demographic information and car ownership/purchasing;
- Applications involving an environmental component
- Applications involving the RDD-based methods
- Applications in ranking/recommending

The last two groups were merged into one for subgroup discussions.

4.2.3 Group notes for applications with an environmental component

This group discussed various cases of spatial causal inference most related to environmental science applications. These applications overall share a similar DAG structure: $Y \leftarrow Z \rightarrow X$ with the potential internal structure of Z where all X , Y , and Z are possibly spatial.

- **quantifying the causal impact of climate variables on arctic sea ice loss**
 - continuous setting with gridded spatial variables
 - dependency structure with spatial interference (?): $X_{s,t-1} \rightarrow Y_{s,t} \leftarrow X_{s+1,t-1}$
 - causal structure at s is conventional 3 variable DAG with confounding
 - question: do all neighbors of s influence Y , or even more (think decaying dependency structure)
 - lagged time dependence, not necessarily Markovian
 - potential spatial dependency structure between the X s
- **Wind farm productivity across wind turbines**
 - Granger causal model
 - wind speed (Y) caused by several variables of wind turbines (humidity, cloud cover, ...)(X)
 - Motivation: when to turn off the turbine
 - Weather forecasts are not sufficient for this as they are not fine-grained enough, and weather models are erroneous
 - What are the functional relationships?
 - Simon: this is more like a predictive problem
 - treat each turbine independently $\rightarrow$ not necessarily spatial
- **Chicago crime counts**
 - Crime scenes in districts in Chicago
 - Can have dependence structure between types of crime or can be spatially correlated crimes (e.g., among neighborhoods)
 - What is the causal relationship?
 - Use causal inference to improve predictive accuracy based on Granger causality
 - Andrew: using a low-rank GP model to predict crimes in space and time is also possible
- **CO2 affects general primary productivity (GPP)**
 - Interested in coefficient: $CO2_t \rightarrow GPP_{st}$
 - distinguish problem into low-frequency and high-frequency components

- Low frequency: $MET_{st} \rightarrow GPP_{st}, S(ize) \rightarrow MET_{st} \rightarrow GPP_{st} \leftarrow S(ize), S(ize) \rightarrow MET_t \rightarrow GPP_t \leftarrow S(ize)$
- High frequency: $CO2_t \rightarrow MET_{st} \rightarrow GPP_{st} \leftarrow CO2_t, N_{st} \rightarrow GPP_{st}, N_{st} \leftarrow S \rightarrow P \rightarrow GPP, CO2_t \leftarrow H \rightarrow N_s$
- fairly complicated DAG structure with supposed hidden confounding
- Question: is the causal effect identifiable? What functional assumptions does one have to make? Can data be pooled or not?
- **Wildfire effects on air quality**
 - $X \leftarrow Z \rightarrow Y \leftarrow X$: similar spatial interference as in Jianwu's study
 - distinguish indirect vs direct treatment
 - paper note: causal inference and wildfires
- **Discussion afterward within the group**
 - What are the peculiarities of the spatial component? So far either not considered or not making an impact on the estimation of causal effects. In these cases, conventional causal/statistical methods should likely be sufficient
 - How do spatial lag and modulators come into the models?

4.2.4 Group notes for applications involving RDD methods

- Thinking about a flow-chart for the specification of a spatial model, possibly based on Akbari's thesis diagram to identify a specification search process to identify specific causal effects [1];
- Where does spatial synthetic control fit here?
- How can you come up with a robustness check procedure to convince others that the identification is effective contextualization of the size of the effect across studies
- Industry is a great example where the size of the effect is quite relevant. Also, it is challenging because the average treatment effect can be quite small, while individual or subgroup treatment effects can be quite large. The spatiality of the predicted individual treatment effects can be useful in visualizing this heterogeneity/uncertainty
- Why not use the higher-level learners (<https://causalml.readthedocs.io/en/latest/methodology.html#t-learner>), rather than the matching, in order to identify the potential high-level spatial confounder? Instead of specifying the covariate relationships by hand, you allow the ML to identify the confounding paths and estimate the nuisance parameter
- It seems that some of the learners do fit into this framework – again possibly a terminological difference/incommensurability. <https://doi.org/10.1080/01621459.2020.1817749> seems to implement at least a T-, possibly X-learner. So, it seems important to provide a map between the kinds of learners to classic identifiability strategies in the PO framing.
- Focusing on the value opt methods might also be useful overlap for the *spatial applications* in planning contexts. In reality, spatial interventions can be quite expensive to implement and/or execute effectively.
- Where does the spatial confounding and/or interference enter this framework? directly into the estimator
- How can you represent the spatial information in the learner framework? The way that proximity is encoded is highly variable – thus how the learner learns the structure can be highly variable. The representation affects the structure with which the data can be introduced into the model. But, this also ensures that you can “reduce” the DAG down?

- What about prior simulation checks? introducing that information just from the prior expected model structure
- Solicit many different judgments about DAGs and then intersect/learn them somehow
- **How to define the exposure mechanism?** Problems with anticipatory effects and “messy” spatial/temporal assignment of treatment. Examples of transit-oriented development where the actual “exposure” to a treatment (building a transit line) or “voting for a winner” effect is a social/non-experimental signal. Can this be represented as a mediator of treatment directly?
- Spatial representation issues in the treatment function, since often, exposure to a spatial treatment is heterogeneous and continuous. **Thinking about ensemble representations of causal effects**, in simulation, you use a very diverse ensemble of potential causal mechanisms, then compare the strength and/or plausibility of each mechanism. One could use non-nested model comparison/averaging across very heterogeneous outcomes, but how is this done in practice? And, just from the difference in the estimation structure, one might expect the estimates to differ – is this true? We could think about generating graphs explicitly to create this ensemble.
- Identifying how there are matches between the larger metalearner frameworks and PO-based investigations. There should be a way for us to specify this as a DAG to help us engage with the spatial vs. nonspatial DAG.

4.2.5 Coming back together from the subgroups

- Discuss the commonalities around representing the spatial confounding
- Relationships hinted with a representational learning problem
- Suggestions that the Airbnb exclusion issue is similar to the one-site feedback problem: a bank wants to know if a person will pay back a loan after receiving one. However, in order to know the outcome, the bank has to give the loan to the person in the first place.
- Presentation of the spatial contextual DAG drafts.

4.3 Day 3: Methodologies Continued

On the morning of day three, we continued the discussions on methodologies of spatial causal inference, following a talk by Kevin Credit on spatial causal forests.

4.3.1 Spatial Causal Model Specifications (Talk by Kevin Credit)

The talk focuses on causal forests from the PO perspective.

4.3.1.1 Spatial autoregressive process

- Spatial autoregressive process (no weighting around treatment variable X)
- SLX (spatial lag of X) (weighting around treatment variable X)
Discussion on the structure of the spatial weights matrix.
- Durbin model (weighting around treatment variable X as well as T)

4.3.1.2 Heterogeneous Treatment Effects (HTE)

The ATE can only be observed across units (“fundamental problem of causal inference”), from the latest Rubin paper (Xie et al 2018). They claim that you cannot learn much from ATE, because of the “averaging” of positive and negative effects. It is therefore useful to look at *heterogeneous* treatment effects.

In **causal forest**: maximise the difference between the TE. Questions:

1. Can we include the treatment in the causal forest?
2. Balance 1's and 0's
3. If you maximise the difference, don't you overestimate the noise, because the goal is to get the biggest split?
4. Are all the covariates known to be pre-treatment? Yes.
5. Table: Training data labeled as T.1 etc. Then check into which leaf the test obs R.1 fall. Then weigh each R.1 by how likely it is to fall into the same leaf as a test observation. Then check which training obs was treated or not, in order to compute the treatment effect on R.1 later.
6. Does it only work for binary treatment? Should work on continuous treatment, need to cycle through all treatments.

4.3.2 Spatial Causal Forest

Start with simulating spatial data (discussion of the code). Spatial causal forest estimates the CATE nearly perfectly.

4.3.3 Afternoon walk

The participants adjourned for an afternoon walk.

4.4 Day 4: Demo of Causal Inference Packages, Benchmarking, and Open Questions

4.4.1 Totte Harinen CausalML demo

- <https://github.com/uber/causalml> CausalML Package on GitHub
- Demonstrates the https://causalml.readthedocs.io/en/latest/examples/uplift_trees_with_synthetic_data.html `UpliftRandomForestClassifier` class
- Also discusses the study on <https://arxiv.org/pdf/2109.05104> estimating the causal effect of personalized climate communication
- Discussion of the gain curve plotted by default with the `UpliftRandomForestClassifier`. Spatially, you may not be able to assume independence of treatment. So, you can't assume that you can separately treat each observation, and the plot of gain as a function of the population quartile isn't valid in this case. You'd need some kind of correction.
- Experimentally, you can introduce continuous levels of treatment
- **How is uncertainty assessed?** Bootstrapping, but this can be tough in big data settings.

4.4.2 Causal Discovery

The goal of causal discovery is to learn causal graphs from data, not necessarily under experiment conditions. This approach relies on observational data when experiments are not sufficient in themselves. Conclusions can be drawn on the graph from which the data has been generated. This is a nontrivial question and the task is unsolvable without assumptions.

Problem needs.

- A clear definition
- A specification of the assumptions being made

Approaches.

- *Constrain-based approach*: causal structure imposes constraints on data from assumptions. Ex: Markov condition, independence, and conditional independence. Usually, this is under-specified: Markov equivalence.
- *Score-based approach*: assign a scoring function that takes data and graph and measures the conformity of the graph to the data, then picks the best scoring graph. The underlying assumption is that the real data-generating process is the minimizer of the score. Define score : $\text{argmin}(S(\text{Graph}, \text{Data}))$ on graph structures.
- *Restricted functional model type approach*: the assumption here is that the data is generated according to a family of models. Fit the model in both directions (X on Y and then Y on X). Valid approach if the fit can only be in the restricted class in one direction. (some just take the one that works better, but this is not necessarily good practice, so perhaps a bad method when one has to decide). Can be rephrased as a score-based approach. Ex: Linear non-gaussian model. LINGAM model: if X and Y are non-gaussian, cannot work in both directions (If it can be done in both directions, then it is Gaussian.)

Discussion. How do you integrate space in the discovery process? If there is background knowledge of any spatial effect then it should be included.

Process of algorithm.

- (i) Specify nodes.
- (ii) Consider the complete graph.
- (iii) Start testing conditional independence between nodes. (the faithfulness assumption is needed here).
- (iv) Remove links until there is no more independence (in the data)
- (v) Then the task is to add the directions: consider all the possibilities

4.4.3 Tigramite Demo

Tigramite (<https://github.com/jakobrunge/tigramite>) is a package for causal inference for time series.

Discover the temporal causal graph using PCMCI (2-step procedure, based on the PC algorithm adding conditional independence on direct parents.):

- Detects the graph and the regime. Effects and causal drivers can change with regime, ex: seasons, turbulence regime, etc., which could be potentially spatial.
- Method with multiple datasets: they could be in different spaces for example. This method discovers the union graph so it pulls the dataset together.

4.4.3.1 Conditional independence tests

- Cases for linearity, nonlinearity (called on by the algorithm)
- These tests can also be applied to vectors (for instance vectors over space), considered equivalent to all the univariate tests (iff vectors are conditionally independent, then the components are also independent).

4.4.3.2 Can a spatial bootstrapping method be included for space independence?

- Can detect when some assumption violation has happened. For instance, the problem of near-deterministic relationships, where variables may need to be grouped in order to measure correlations.
- *What happens to the graph when the data is grouped in a certain way? (See tutorial)*

4.4.3.3 Case Study: Circulation of air mass in different regions

- Finding the causal graph where nodes are circulation in different areas, or estimating the causal effects based on a DAG by experts.
- *Stationarity assumption needed* This can fix part of the graph for all these methods.
- In the situation with multiple datasets, the assumption is that these datasets have been generated independently

Note: *Question of intervention in time series is a tricky one, what exactly is an intervention? How about a spatial intervention?*

4.4.4 Benchmarking and Evaluation

4.4.4.1 Causal discovery evaluation without ground truth by Dominik Janzing

Dominik Janzing presents the study on causal discovery benchmarking without ground truth [3]

- Deterministic relationship vs noisy relationship
- Check causality falsifiable
- If $X- > Y$ in the unconfounded setting and $Y- > Z$ in the unconfounded setting, we only have $X- > Y- > Z$, no $X- > Z$.
- Unconfounded: $P(Z|do(Y)) = P(Z|Y)$
- With generalization assumption: X independentOf $Z|Y$
- We could have independent X_1 and X_2 within X that $X_1- > Y$, $X_2- > Z$ and $Y- > Z$.
- Pick sufficient subsets from the original set, and apply causal discovery algorithms to the subsets. Then check the compatibility of the results.
- Compatibility: A is a causal discovery algorithm, $A(X_1, X_2, \dots, X_n)$ and $A(X_2, \dots, X_n, X_{n+1})$
- Graph 1: $X_1- > X_2- > X_3$ and $X_1- > X_3$, $P(X_3|do(X_2)) <> P(X_3|X_2)$
- Graph 2: $X_2- > X$, $P(X_3|do(X_2 = x_2)) = P(X_3|X_2)$
- ADMG: Acyclic Directed Mixed Graph
- Causality is to generalizing unseen variables
- The idea is similar to bootstrapping (leaving some variables)
- PC depends on the ordering of the variables. There are efforts to make the ordering irrelevant.
- Compatibility with data: predict unseen variables

4.4.4.2 Spatial causal inference synthetic data shown by Jianwu Wang

Jianwu Wang presents the case study [2], (Pre-Print: <https://arxiv.org/abs/2405.08174>)

- **What is the ground truth?** Will be the coefficient for the linear model.
- Noise can be added
- More general versions: 1) spatial processes, 2) distance function, 3) adjacency matrices, 4) different agents interact with each other spatially
- Frequency domain could help identify independent causal graphs.
- Intervention is one-time or lasting, continuous/binary.
- Noted that Jacob Runge's group is also working on benchmarking causal models, particularly causal discovery models

4.4.5 Open Questions

In this session, seminar participants identified key open questions on the topic of spatial causal inference and indicated interest in follow-up discussion and research (in [] below).

1. Spatio-temporal extension of Granger causality to 2+ variables [Katerina]
2. Modelling spatial neighborhood as a moderator in a causal process [Katerina]
3. Investigation of marginal conditions/data limits on causal claims (spatial area imposes discontinuity on the process, temporal bounds limit the ability to capture the entire process, resolution matters [Martin]
4. Chain graphs and representations for spatial characteristics in DAGs (incl. nuancing for CAR and SAR processes) [Levi, Martin, Andreas, Cecile, Jonas, Yanan]
5. Robustness of spatial causal claims to noise in observations, and to sampling granularity. Impact of uncertainty in observed outcome variables and independent variables on causal claims and discovery. [Martin, Levi]
6. Parameter identifiability in process models. Can different parameters lead to the same observed outcome distributions? [Jonas, Andreas] [this may also map to 5]
7. Equifinality – can the same DAG lead to different interventional results? [Cecile] [maps to 6?]
8. Interventions w spatial targeting, including experimental study with regional targeting (Uber) [Totte, Levi, Martin, possibly Kevin]
9. How to quantify Treatment effect from Point process on point process. [Shu, Cecile, Andrew, Katerina]
10. **Adding spatial statistics to causality/reframing spatial causality via spatial statistics.** Relates/informs [5, 2, 1, 12] – Main conceptual paper [All. MT coordinates] [whiteboard included Andrew, Shu, Urmi, Levi, MT, Katerina.]
11. Shared benchmark task and resource [Jianwu, Jonas, all]. links to [4], could link to 10.
12. **Whitepaper** – Spatial Causal Inference Framework – starts as position paper [10], informs all other papers, and provides conceptual/philosophical grounding. [all]
13. Causal inference on spatial networks [Martin]

4.5 Day 5: Revisiting the Definition of Spatial Causal Inference

4.5.1 Back to the beginning: Towards a definition of spatial causal inference

In this session we returned to the discussion from Day 1, revisiting the definition of *spatial* causal inference.

Spatial causal inference is the case of causal inference where the spatial context of the process (or its parts, i.e., differentiable in space) matters.

In other words, spatial problems are those where map randomisation (i.e., permutation of the spatial encoding of the variables, such as shuffling geometries, or grid indices) breaks spatial associations.

Testing whether the causal relationships are independent of space is done through *negation*, assuming a phenomenon *is* spatial, unless:

- **Process perspective:** (affine) spatially continuous transformations of the spatial process lead to invariant outcomes. If true, the results are invariant.
- **Spatial permutation test (for discrete data):** spatial confounding variables are shuffled to check whether the association with outcomes remains invariant. If true, this is not a spatial problem. Shuffling may include permutations of attribute values of the spatial features (i.e., substitution of values to the geometries) or of spatial indices of grids.
- **Altering imposed aggregation boundaries:** enacting MAUP, altered spatial aggregations may break the causal relationships.
- space subset of the variables of a spatial process depends on space, we call it a spatial process
- question: will the permutation change the *do* operation

Caution: Not all processes transformations/data permutations/aggregations break the associations! The robustness of these tests needs to be carefully considered.

In this perspective, one begins by assuming space matters in the causal relationships, and evaluates the *joint spatial distribution* of independent and dependent variables. Both independent variables, as well as outcome (dependent) variables, may vary in space. A special case is where none of these vary in space (the processes are stationary), but their association is not stationary. For example, consider a stationary process, where the outcomes are a spatially translated version of this process, e.g., a spatially uncorrelated process occurring on a tectonic plate shifting.

4.5.2 Possible tests

- A permutation test similar to the Moran's I test

4.5.3 Spatial causal inference and its link to causal inference

Spatiotemporal causal inference *covers* the simplified case of causal inference, and tests must be made for the special case where space may be ignored. We need to link the above spatial characteristics to causal inference, and the definitions by Pearl, Rubin, and Dominik.

4.5.4 Position paper

The group intends to write a position paper on the above for a potential venue such as PNAS, or Philosophical Transactions of the Royal Society.

4.5.5 Following up workshops and activities

The group intends to organize related workshops or seminars to maintain the momentum of the exchange on spatial causal inference. Some potential venues include ACM SIGSPATIAL workshops, GIScience workshops, Dagstuhl Seminars, NeurIPS, etc. During the writing of this report, the 1st ACM SIGSPATIAL International Workshop on Spatiotemporal Causal Analysis (STCausal 2024) has been accepted and will take place on Oct. 29th, 2024.

Further, a central web domain www.spatial-causal.org has been registered to act as a center point for the activities spilling from this seminar in the future (currently pointing to the 2024 seminar).

References

- 1 Kamal Akbari, Stephan Winter, and Martin Tomko. Spatial causality: A systematic review on spatial causal inference. *Geographical Analysis*, 55(1):56–89, 2023.
- 2 Sahara Ali, Uzma Hasan, Xingyan Li, Omar Faruque, Akila Sampath, Yiyi Huang, Md Osman Gani, and Jianwu Wang. Causality for earth science—a review on time-series and spatiotemporal causality methods. *arXiv preprint arXiv:2404.05746*, 2024.
- 3 Philipp M. Faller, Leena Chennuru Vankadara, Atalanti A. Mastakouri, Francesco Locatello, and Dominik Janzing. Self-compatibility: Evaluating causal discovery without ground truth. AISTATS 2024.

Participants

- Kevin Credit
Maynooth University, IE
- Cécile de Bézenac
Alan Turing Institute – London,
GB & University of Leeds, GB
- Simon Dirmeier
Swiss Data Science Center –
Zürich, CH
- Andreas Gerhardus
DLR – Jena, DE
- Totte Harinen
Airbnb – San Francisco, US
- Dominik Janzing
Amazon Web Services –
Tübingen, DE
- Urmi Ninad
TU Berlin, DE
- Markus Reichstein
MPI für Biogeochemie –
Jena, DE
- Katerina Schindlerova
Universität Wien, AT
- Martin Tomko
University of Melbourne –
Carlton, AU
- Jonas Wahl
TU Berlin, DE
- Jianwu Wang
University of Maryland –
Baltimore County, US
- Levi John Wolf
University of Bristol, GB
- Yanan Xin
ETH Zürich, CH
- Shu Yang
North Carolina State University –
Raleigh, US
- Andrew Zammit Mangion
University of Wollongong, AU

